

REMARKS

In response to the Office Action mailed January 30, 2004, Applicant respectfully requests reconsideration. To further the prosecution of this Application, Applicant submits the following remarks and has added new claims. The claims as now presented are believed to be in allowable condition.

Claims 1-22 were pending in this Application. By this Amendment, claims 23-30 have been added. Accordingly, claims 1-30 are now pending in this Application. Claims 1, 7, 14 and 21 are independent claims.

Objection to the Specification

The Abstract of the Specification was objected to due to a minor informality. Applicant has amended the Specification to cure this minor informality. Accordingly, the objection to the Specification should be withdrawn.

Rejections under §103

Claims 1-12, 14-19 and 21-22 were rejected under 35 U.S.C. §103(a) as being unpatentable over U.S. Patent No. 5,884,098 (Mason, Jr.). Claims 13 and 20 were rejected under 35 U.S.C. §103(a) as being unpatentable over Mason, Jr. in view of European Patent 0114190 A2 (Hartung).

Applicant respectfully traverses each of these rejections and requests reconsideration. The claims are in allowable condition because that patentably distinguish over the cited prior art.

Mason, Jr. discloses a RAID disk controller 201 having a front end cache 205 and a back end cache 209 (column 5, lines 6-24 and Fig. 2). The front end cache 205 and the back end cache 209 are separate software processes executing on one or more microcontrollers which exercise control over all of the disk controller 201 (column 5, lines 24-36). Although only one cache memory is used, the front end cache 205 and the back end cache 209 may either share a single control store or use separate control stores to hold control structures such as a least recently used (LRU) block queue (column 6, lines 7-11). For

convenience, the following description refers to the queue of blocks available in each cache system simply as a front end cache block list and back end cache block list (column 6, lines 11-13). LRU queue algorithms and techniques are well known (column 6, line 14). In particular, an LRU cache block list may be a data structure configured as a doubly linked list holding a pointer to each cache block currently allocated in the cache memory (column 6, lines 14-18). When a block which is already in the queue is used, the entry for that block in the list is moved to the head of the list (column 6, lines 18-20). The entry in the list corresponding to the block which is used at the earliest time (or even never used) eventually moves to the last position in the list (column 6, lines 20-23). When the cache memory becomes full, subsequent attempts to allocate cache blocks cause the last entry in the list to be removed, deallocating the cache block pointed to (column 6, lines 23-26). The cache memory block previously pointed to by the last entry in the list is then replaced with the new block for which allocation was attempted and an entry for the new block placed at the head of the list (column 6, lines 26-29). To optimize the performance of this structure, it is desirable to coordinate the operations performed by the front end cache 205 and the back end cache 209 (column 6, lines 30-33). In particular, there is a communication path 211 established between the front end cache 205 and the back end cache 209 (column 6, lines 52-55). The communication path 211 established between the front end cache 205 and the back end cache 209 may take several forms including communications hardware built into the disk array controller 201 (e.g., one microprocessor controlled by a second microprocessor over a dedicated serial line, two processors merged into a single hardware entity performing both functions in response to a multitasking software system, a system bus or shared memory (column 6, line 60 through column 7, line 12). The presently preferred communication path 211 is a common data structure residing in a common control store accessible to both the front end cache 205 and the back end cache 209 (column 7, lines 13-16).

Hartung discloses methods and apparatus for peripheral data handling hierarchies (Title). In connection with data processing systems manufactured by IBM, there is a term "channel commands" (page 12, lines 1-3). A user identification (USEID) accompanies the ESTABLISH REFERENCE CHARACTERISTICS (ERC) command (page 12, lines 6-15). When erasure of USEIDs indicate detection of a control error, then a DCB 61 USEID section is examined (page 12, lines 21-36). Additionally, a parity error indicates an error in control data structures, and other error detection codes may also be employed (page 26, lines 4-5).

#### Claims 1-6

Claim 1 is directed to a memory board for a data storage system. The memory board includes an interface which is configured to couple to a bus of the data storage system, memory which is configured to store a doubly linked list data structure, and a memory board control circuit coupled to the interface and the memory. The memory board control circuit is configured to receive a modify command from a processor of a data storage system through the interface. The processor is configured to move data within the data storage system. The memory board control circuit is further configured to atomically modify the doubly linked list data structure in accordance with the modify command, and provide a result to the processor of the data storage system through the interface in response to modifying the doubly linked list data structure.

The cited prior art does not teach or suggest a memory board for a data storage system which includes a memory board control circuit where the memory board control circuit is configured to atomically modify the doubly linked list data structure in accordance with a modify command, as recited in claim 1. Rather, Mason, Jr. discloses a RAID disk controller 201 having a front end cache 205 and a back end cache 209, and a communication path 211 which is used to coordinate operations performed by the front end cache 205 and the back end cache 209 to optimize the performance of a least recently used cache block

(LRU) configured as a doubly-linked list (e.g., see column 6, line 7 through column 7, line 16 of Mason, Jr.).

It should be understood that the Mason, Jr. disk controller 201 is similar to Applicant's admitted prior art in which multiple CPU boards employ data structure sharing techniques to prevent the CPU boards from inadvertently accessing the LRU queue simultaneously (see page 3, line 25 through page 4, line 1 of the Specification). In Mason, Jr., the front end cache 205 and the back end cache 209 coordinate their operations by communicating through a communication path 211 which is, for example, a common data structure residing in a common control store accessible to both the front end cache 205 and the back end cache 209 (e.g., see column 7, lines 13-16 of Mason, Jr.). A benefit of the invention as recited in claim 1 is that atomically modifying a doubly linked list data structure alleviates the need for shared data structure locking overhead as explained in the Specification, for example, on page 5, line 24 through page 6, lines 6 of the Specification. There is no teaching or suggestion of atomically modifying a doubly linked list data structure, as recited in claim 1.

Furthermore, it is unclear how one could modify the Mason, Jr. disk controller 201 to atomically modify a doubly linked list data structure or why one would want to modify the Mason, Jr. disk controller 201 since the Mason, Jr. disk controller 201 is already equipped with a communication path 211 to coordinate operations. Hartung, which was cited against claims 13 and 20 as disclosing error checking, does not teach or suggest how one could modify the Mason, Jr. disk controller 201 in that manner or why one would want to modify the Mason, Jr. disk controller 201 in that manner.

For the reasons stated above, claim 1 patentably distinguishes over the cited prior art, and the rejection of claim 1 under 35 U.S.C. §103(a) should be withdrawn. Accordingly, claim 1 is in allowable condition.

Because claims 2-6 depend from and further limit claim 1, claims 2-6 are in allowable condition for at least the same reasons. Additionally, it should be

understood that the dependent claims recite additional features which further patentably distinguish over the cited prior art.

For example, claim 6 further recites limitations of (i) the memory board control circuit being configured to atomically modify the doubly linked list shared data structure by performing a series of individual transactions on the doubly linked list shared data structure, and (ii) the memory board control circuit being configured to provide, as a result, a series of transaction outputs respectively corresponding to the series of individual transactions. There is no such teaching or suggestion in the cited prior art. In particular, there is no mention of atomic modification in Mason, Jr. Furthermore, there is no mention of atomic modification or for that matter a memory board control circuit being configured to provide, as a result, a series of transaction outputs respectively corresponding to the series of individual transactions in Hartung. The Office Action contends on page 23, last paragraph, that Hartung discloses parity error checking on the quality of data transfer. Regardless of whether this contention is correct, there is no disclosure in Hartung of a memory board control circuit being configured to atomically modify the doubly linked list shared data structure by performing a series of individual transactions on the doubly linked list shared data structure, and further being configured to provide, as a result, a series of transaction outputs respectively corresponding to the series of individual transactions, as recited in claim 6. If the rejection of claim 6 is be maintained in view of Hartung, Applicant respectfully requests that it be pointed out with particularity where Hartung provides such a disclosure.

### Claims 7-13

Claim 7 is directed to a data storage system having a set of storage devices, a processor which is configured to move data to and from the set of storage devices, a bus coupled to the processor, and a memory board. The memory board includes (i) an interface which couples to the bus, (ii) memory which is configured to store a doubly linked list data structure, and (iii) a memory

board control circuit coupled to the interface and the memory. The memory board control circuit is configured to receive a modify command from the processor of the data storage system through the interface and the bus, atomically modify the doubly linked list data structure in accordance with the modify command, and provide a result to the processor of the data storage system through the interface and the bus in response to modifying the doubly linked list data structure.

As explained above in connection with claim 1, the cited prior art does not teach or suggest such a memory board control circuit. Accordingly, the cited prior art also does not teach a data storage system having such a memory board control circuit. As a result, claim 7 patentably distinguishes over the cited prior art for at least the same reasons as claim 1. Therefore, the rejection of claim 7 under 35 U.S.C. §103(a) should be withdrawn, and claim 7 is in allowable condition.

Because claims 8-13 depend from and further limit claim 7, claims 8-13 are in allowable condition for at least the same reasons.

#### Claims 14-20

Claim 14 is directed to a method for accessing a doubly linked list data structure. The method is performed in a memory board of a data storage system and includes the step of receiving a modify command from a processor of a data storage system through a bus of the data storage system. The processor is configured to move data within the data storage system. The method further includes the steps of atomically modifying the doubly linked list data structure in accordance with the modify command, and providing a result to the processor of the data storage system through the bus in response to modifying the doubly linked list data structure.

As explained above in connection with claim 1, the cited prior art does not teach or suggest such a method because the cited prior art does not disclose atomically modifying a doubly linked list data structure in accordance with a

modify command. Accordingly, the cited prior art also does not teach a method having such a step. As a result, claim 14 patentably distinguishes over the cited prior art for at least the same reasons as claim 1. Thus, the rejection of claim 14 under 35 U.S.C. §103(a) should be withdrawn, and claim 14 is in allowable condition.

Because claims 15-20 depend from and further limit claim 14, claims 15-20 are in allowable condition for at least the same reasons.

#### Claims 21-22

Claim 21 is directed to a memory board control circuit for accessing a doubly linked list data structure of a data storage system. The memory board control circuit is mountable to a memory board. The memory board control circuit includes an input port that couples to a bus of the data storage system, an output port that couples to the bus of the data storage system, and control logic connected to the input port and to the output port. The control logic is configured to receive a modify command from a processor of a data storage system through the input port. The processor is configured to move data within the data storage system. The control logic is further configured to atomically modify the doubly linked list data structure in accordance with the modify command, and provide a result to the processor of the data storage system through the output port in response to modifying the doubly linked list data structure.

As explained above in connection with claim 1, the cited prior art does not teach or suggest a memory board control circuit configured to atomically modify a doubly linked list data structure in accordance with a modify command. Accordingly, claim 21 patentably distinguishes over the cited prior art for at least the same reasons as claim 1. Thus, the rejection of claim 21 under 35 U.S.C. §103(a) should be withdrawn, and claim 21 is in allowable condition.

Because claim 22 depends from and further limits claim 21, claim 22 is in allowable condition for at least the same reasons.

Newly Added Claims

Claims 23-30 have been added and are believed to be in allowable condition. Claims 23-24 depend from claim 1. Claims 25-26 depend from claim 7. Claims 27-28 depend from claim 14. Claims 29-30 depend from claim 20. Support for claims 23-30 is provided within the Specification, for example, on page 5, line 24 through page 6, line 6; page 8, line 22 through page 9, line 6; page 21, lines 17-29; and Fig. 4. No new matter has been added.

Conclusion

In view of the foregoing remarks, this Application should be in condition for allowance. A Notice to this affect is respectfully requested. If the Examiner believes, after this Amendment, that the Application is not in condition for allowance, the Examiner is respectfully requested to call the Applicant's Representative at the number below.

Applicant hereby petitions for any extension of time which is required to maintain the pendency of this case. If there is a fee occasioned by this Amendment, including an extension fee, that is not covered by an enclosed check, please charge any deficiency to Deposit Account No. 50-0901.

-20-

If the enclosed papers or fees are considered incomplete, the Patent Office is respectfully requested to contact the undersigned collect at (508) 366-9600, in Westborough, Massachusetts.

Respectfully submitted,



David E. Huang, Esq.  
Attorney for Applicant  
Registration No.: 39,229  
CHAPIN & HUANG, L.L.C.  
Westborough Office Park  
1700 West Park Drive  
Westborough, Massachusetts 01581  
Telephone: (508) 366-9600  
Facsimile: (508) 616-9805

Attorney Docket No.: EMC01-19(01056)

Dated: April 30, 2004